

Transformation Forests

Torsten Hothorn
Universität Zürich

Achim Zeileis
Universität Innsbruck

Abstract

Regression models for supervised learning problems with a continuous target are commonly understood as models for the conditional mean of the target given predictors. This notion is simple and therefore appealing for interpretation and visualisation. Information about the whole underlying conditional distribution is, however, not available from these models. A more general understanding of regression models as models for conditional distributions allows much broader inference from such models, for example the computation of prediction intervals. Several random forest-type algorithms aim at estimating conditional distributions, most prominently quantile regression forests (Meinshausen, 2006, JMLR). We propose a novel approach based on a parametric family of distributions characterised by their transformation function. A dedicated novel “transformation tree” algorithm able to detect distributional changes is developed. Based on these transformation trees, we introduce “transformation forests” as an adaptive local likelihood estimator of conditional distribution functions. The resulting models are fully parametric yet very general and allow broad inference procedures, such as the model-based bootstrap, to be applied in a straightforward way.

Keywords: random forest, transformation model, quantile regression forest, conditional distribution, conditional quantiles.

1. Introduction

Supervised machine learning plays an important role in many prediction problems. Based on a learning sample consisting of N pairs of target value y and predictors \mathbf{x} , one learns a rule r that predicts the status of some unseen Y via $r(\mathbf{x})$ when only information about \mathbf{x} is available. Both the machine learning and statistics communities differentiate between “classification problems”, where the target Y is a class label, and “regression problems” with conceptually continuous target observations y . In binary classification problems with $Y \in \{0, 1\}$ the focus is on rules r for the conditional probability of Y being 1 given \mathbf{x} , more formally $\mathbb{P}(Y = 1 \mid \mathbf{X} = \mathbf{x}) = r(\mathbf{x})$. Such a classification rule r is probabilistic in the sense that one cannot only predict the most probable class label but also assess the corresponding probability. This additional information is extremely valuable because it allows an assessment of the rules’ r uncertainty about its prediction. It is much harder to obtain such an assessment of uncertainty from most contemporary regression models, because the rule (or “regression function”) r typically describes the conditional expectation $\mathbb{E}(Y \mid \mathbf{X} = \mathbf{x}) = r(\mathbf{x})$ but not the full predictive distribution of Y given \mathbf{x} . Thus, the prediction $r(\mathbf{x})$ only contributes information about the mean of some unseen target Y but tells us nothing about other characteristics of its distribution. Without making additional restrictive assumptions, for example

constant variances in normal distributions, the derivation of probabilistic statements from the regression function r alone is impossible.

Contemporary random forest-type algorithms also strongly rely on the notion of regression functions r describing the conditional mean $\mathbb{E}(Y \mid \mathbf{X} = \mathbf{x})$ only (for example [Biau *et al.* 2008](#); [Biau 2012](#); [Scornet *et al.* 2015](#)), although the first random forest-type algorithm for the estimation of conditional distribution functions was published more than a decade ago (“bagging survival trees”, [Hothorn *et al.* 2004](#)). Similar approaches were later developed independently by [Lin and Jeon \(2006\)](#), using the name “adaptive nearest neighbours”, by [Meinshausen \(2006\)](#) in his “quantile regression forests” and by [Ishwaran *et al.* \(2008\)](#) as “random survival forests”. All these random forest-type algorithms are based on “nearest neighbour weights”. The idea is to obtain a “distance” measure based on the number of times a pair of observations is assigned to the same terminal node in the different trees of the forest. Similar observations have a high probability of ending up in the same terminal node whereas this probability is low for quite different observations. Then, the prediction for predictor values \mathbf{x} (either new or observed) is simply obtained as a weighted empirical distribution function (or Kaplan-Meier estimator in the context of right-censored target values) where those observations from the learning sample similar (or unsimilar) to \mathbf{x} in the forest receive high (or low/zero) weights, respectively. Although this aggregation procedure in the aforementioned algorithms is suitable for estimating predictive distributions, the underlying trees are not. The reason is that the ANOVA- or log-rank-type split procedures commonly applied are not able to deal with distributions in a general sense. Consequently, the splits favour the detection of changes in the mean – or have power against proportional hazards alternatives in survival trees. However, in general, they have very low power for detecting other patterns of heterogeneity (e.g., changes in variance) even if these can be explained by the predictor variables. A simple toy example illustrating this problem is given in Figure 1. Here, the target’s conditional normal distribution has a variance split at value 0.5 of a uniform $[0, 1]$ predictor. We fitted a quantile regression forest ([Meinshausen 2006, 2016](#)) to the 10,000 observations depicted in the figure along with ten additional independent uniformly distributed non-informative predictors (using 100 trees, three selected variables in each node and a minimum nodesize of 50 observations in the forest). The true conditional 20% and 80% quantiles are not approximated very well by the quantile regression forest. In particular, the split at 0.5 does not play an important role in this model. Thus, although such an abrupt change in the distribution can be represented by a binary tree, the traditional ANOVA split criterion employed here was not able to detect this split.

[Figure 1 about here.]

To improve upon quantile regression forests and similar procedures in situations where changes in moments beyond the mean are important, we propose “transformation forests” for the estimation and prediction of conditional distributions for Y given predictor variables \mathbf{x} and proceed in three steps. We first suggest to understand forests as adaptive local likelihood estimators. Second, we recap the most important features of the flexible and computational attractive “transformation family” of distributions ([Hothorn *et al.* 2014, 2016](#)) which includes a variety of distribution families. Finally, we adapt the core ideas of “model-based recursive partitioning” ([Zeileis *et al.* 2008](#)) to this transformation family and introduce novel algorithms for “transformation trees” and “transformation forests” for the estimation of conditional distribution functions which potentially vary in the mean and also in higher moments as a

function of predictor variables \mathbf{x} . In our small example in Figure 1, these novel transformation forests were able to recover the true conditional distributions much more precisely than quantile regression forests.

Owing to the fully parametric nature of the predictive distributions that can be obtained from these novel methods, model inference procedures, such as variable importances, independence tests or model-based resampling, can be formulated in a very general and straightforward way. The performance of unconditional and linear transformation models, of quantile regression forests as well as transformation trees and transformation forests was evaluated empirically on the Abalone, Boston Housing, BigMac, and Ozone regression problems as well as artificial variations thereof.

2. Adaptive Local Likelihood Trees and Forests

We first deal with the unconditional distribution \mathbb{P}_Y of a random variable $Y \in \mathcal{Y}$ and we restrict our attention to a specific probability model defined by the parametric family of distributions

$$\mathbb{P}_{Y,\Theta} = \{\mathbb{P}_{Y,\boldsymbol{\vartheta}} \mid \boldsymbol{\vartheta} \in \Theta\}$$

with parameters $\boldsymbol{\vartheta}$ and parameter space $\Theta \subseteq \mathbb{R}^P$. With predictors $\mathbf{X} \in \mathcal{X}$ from some predictor sample space \mathcal{X} , our main interest is in the conditional probability $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}}$ and we assume that this conditional distribution is a member of the family of distributions introduced above, i.e., we assume that a parameter $\boldsymbol{\vartheta}(\mathbf{x}) \in \Theta$ exists such that $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}} = \mathbb{P}_{Y,\boldsymbol{\vartheta}(\mathbf{x})}$. We call $\boldsymbol{\vartheta} : \mathcal{X} \rightarrow \Theta$ the “conditional parameter function” and the task of estimating the conditional distributions $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}}$ for all \mathbf{x} reduces to the problem of estimating this conditional parameter function.

From the probability model $\mathbb{P}_{Y,\Theta}$ we can derive the log-likelihood contribution $\ell_i : \Theta \rightarrow \mathbb{R}$ for each of N independent observations $(y_i, \mathbf{x}_i), i = 1, \dots, N$. We propose and study a novel random forest-type estimator $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N$ of the conditional parameter function $\boldsymbol{\vartheta}$ in the class of adaptive local likelihood estimators of the form

$$\hat{\boldsymbol{\vartheta}}^N(\mathbf{x}) := \arg \max_{\boldsymbol{\vartheta} \in \Theta} \sum_{i=1}^N w_i^N(\mathbf{x}) \ell_i(\boldsymbol{\vartheta}); \quad \mathbf{x} \in \mathcal{X} \quad (1)$$

where $w_i^N : \mathcal{X} \rightarrow \mathbb{R}^+$ is the “conditional weight function” for observation i given a specific configuration \mathbf{x} of the predictor variables (which may correspond to an observation from the learning sample or to new data). This weight measures the similarity of the two distributions $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}_i}$ and $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}}$ under the probability model $\mathbb{P}_{Y,\Theta}$. The main idea is to obtain a large weight for observations i which are “close” to \mathbf{x} in light of the model and essentially zero in the opposite case. The superscript N indicates that the weight function may depend on the learning sample, and in fact the choice of the weight function w_i^N is crucial in what follows.

Local likelihood estimation goes back to Brillinger (1977) in a comment to Stone (1977) and was the topic of Robert Tibshirani’s PhD thesis, published in Tibshirani and Hastie (1987). Early regression models in this class were based on the idea of fitting polynomial models locally within a fixed smoothing window. Adaptivity of the weights refers to an \mathbf{x} -dependent, non-constant smoothing window, i.e., different weighing schemes are applied in different parts

of the predictor sample space \mathcal{X} . An overview on local likelihood procedures was published by Loader (1999). Subsequently, we illustrate how classical maximum likelihood estimators, model-based trees, and model-based forests can be embedded in this general framework by choosing suitable conditional weight functions and plugging these into (1).

The unconditional maximum likelihood estimator is based on unit weights $w_{\text{ML},i}^N := 1$ not depending on \mathbf{x} , i.e., all observations in the learning sample are considered to be equally “close”. And thus

$$\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N := \arg \max_{\boldsymbol{\vartheta} \in \Theta} \sum_{i=1}^N \ell_i(\boldsymbol{\vartheta}).$$

In contrast, model-based trees can adapt to the learning sample by employing rectangular splits to define a partition $\mathcal{X} = \dot{\bigcup}_{b=1, \dots, B} \mathcal{B}_b$ of the predictor sample space. Each of the B cells then contains a different local unconditional model. More precisely, the conditional weight function $w_{\text{Tree},i}^N$ is simply an indicator for \mathbf{x}_i and \mathbf{x} being elements of the same terminal node so that only observations in the same terminal node are considered to be “close”. The weight and parameter functions are

$$\begin{aligned} w_{\text{Tree},i}^N(\mathbf{x}) &:= \sum_{b=1}^B I(\mathbf{x} \in \mathcal{B}_b \wedge \mathbf{x}_i \in \mathcal{B}_b) \\ \hat{\boldsymbol{\vartheta}}_{\text{Tree}}^N(\mathbf{x}) &:= \arg \max_{\boldsymbol{\vartheta} \in \Theta} \sum_{i=1}^N w_{\text{Tree},i}^N(\mathbf{x}) \ell_i(\boldsymbol{\vartheta}). \end{aligned} \tag{2}$$

Thus, this essentially just picks the parameter estimate from the b -th terminal node which is associated with cell \mathcal{B}_b

$$\hat{\boldsymbol{\vartheta}}_b^N = \arg \max_{\boldsymbol{\vartheta} \in \Theta} \sum_{i=1}^N I(\mathbf{x}_i \in \mathcal{B}_b) \ell_i(\boldsymbol{\vartheta}),$$

along with the corresponding conditional distribution $\mathbb{P}_{Y, \hat{\boldsymbol{\vartheta}}_b^N}$. Model-based recursive partitioning (MOB, Zeileis *et al.* 2008) is one representative of such a tree-structured approach.

A forest of T trees is associated with partitions $\mathcal{X} = \dot{\bigcup}_{b=1, \dots, B_t} \mathcal{B}_{tb}$ for $t = 1, \dots, T$. The b -th terminal node of the t -th tree contains the parameter estimate $\hat{\boldsymbol{\vartheta}}_{tb}^N$ and the t -th tree defines the conditional parameter function $\hat{\boldsymbol{\vartheta}}_{\text{Tree},t}^N(\mathbf{x})$. We define the forest conditional parameter function via “nearest neighbour” (using the term coined by Lin and Jeon 2006) forest weights

$$\begin{aligned} w_{\text{Forest},i}^N(\mathbf{x}) &:= \sum_{t=1}^T \sum_{b=1}^{B_t} I(\mathbf{x} \in \mathcal{B}_{tb} \wedge \mathbf{x}_i \in \mathcal{B}_{tb}) \\ \hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x}) &:= \arg \max_{\boldsymbol{\vartheta} \in \Theta} \sum_{i=1}^N w_{\text{Forest},i}^N(\mathbf{x}) \ell_i(\boldsymbol{\vartheta}). \end{aligned} \tag{3}$$

The conditional weight function $w_{\text{Forest},i}^N$ counts how many times \mathbf{x}_i and \mathbf{x} are element of the same terminal node in each of the T trees, i.e., captures how “close” the observations are on average across the trees in the forest. Hothorn *et al.* (2004) first suggested these weights for

the aggregation of T survival trees. The same weights have later been used by Lin and Jeon (2006), Meinshausen (2006) and Ishwaran *et al.* (2008). An “out-of-bag” version only counts the contribution of the t -th tree for observation i when i was not used for fitting the t -th tree. Forests relying on the aggregation scheme (3) model the conditional distribution $Y \mid \mathbf{X} = \mathbf{x}$ for some configuration \mathbf{x} of the predictors as $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})} \in \mathbb{P}_{Y, \Theta}$. In this sense, such a forest is a fully specified parametric model with (in-bag or out-of-bag) log-likelihood

$$\sum_{i=1}^N \ell_i \left(\hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x}_i) \right)$$

allowing a broad range of model inference procedures to be directly applied as discussed in Section 5. Although this core idea seems straightforward to implement, we unfortunately cannot pick our favourite tree-growing algorithm and mix it with some parametric model as two critical problems remain to be addressed in this paper. First, none of the standard tree-growing algorithms is ready to be used for finding the underlying partitions because their variable and split selection procedures have poor power for detecting distributional changes which are not linked to changes in the mean as was illustrated by the simple toy example presented in the introduction. A tailored tree-growing algorithm inspired by model-based recursive partitioning also able to detect changes in higher moments is introduced in Section 4. The second problem is associated with the parametric families $\mathbb{P}_{Y, \Theta}$. Although, in principle, all classical probability models are suitable in this general framework, different parameterisations render unified presentation and especially implementation burdensome. We address this second problem by restricting our implementation to a novel transformation family of distributions. Theoretically, this family contains all univariate probability distributions \mathbb{P}_Y and practically close approximations thereof. We highlight important aspects of this family and the corresponding likelihood function in the next section.

3. Transformation Models

A transformation model $\mathbb{P}(Y \leq y) = F_Y(y) = F_Z(h(y))$ describes the distribution function of Y by an unknown monotone increasing transformation function h and some a priori chosen continuous distribution function F_Z . We use this framework because simple, e.g., linear, transformation functions implement many of the classical parametric models whereas more complex transformation functions provide similar flexibility as models from the non-parametric world. In addition, discrete and continuous targets, also under all forms of random censoring and truncation, are handled in a unified way. As a consequence, our corresponding “transformation forests” will be applicable to a wide range of targets (discrete, continuous with or without censoring and truncation, counts, survival times) with the option to gradually move from simple to very flexible models for the conditional distribution functions $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})}$.

In more detail, let $Z \sim \mathbb{P}_Z$ denote an absolutely continuous random variable with density, distribution and quantile functions f_Z , F_Z and F_Z^{-1} , respectively. We furthermore assume $0 < f_Z(z) < \infty \forall z \in \mathbb{R}$, $F_Z(-\infty) = 0$ and $F_Z(\infty) = 1$ for a log-concave density f_Z as well as the existence of the first two derivatives of the density $f_Z(z)$ with respect to z , both derivatives shall be bounded. We do not allow any unknown parameters for this distribution. Possible choices include the standard normal, the standard logistic and the minimum extreme value

distribution with distribution functions $F_Z(z) = \Phi(z)$, $F_Z(z) = F_{\text{SL}}(z) = (1 + \exp(-z))^{-1}$ and $F_Z(z) = F_{\text{MEV}}(z) = 1 - \exp(-\exp(z))$, respectively.

Let $\mathcal{H} = \{h : \mathcal{Y} \rightarrow \mathbb{R} \mid h(y_1) < h(y_2) \forall y_1 < y_2 \in \mathcal{Y}\}$ denote the space of all strictly monotone transformation functions. With the transformation function h we can write F_Y as $F_Y(y \mid h) = F_Z(h(y)) \forall y \in \mathcal{Y}$ with density $f_Y(y \mid h)$ and there exists a unique transformation function $h = F_Z^{-1} \circ F_Y$ for all distribution functions F_Y (Hothorn *et al.* 2016). A convenient feature of characterising the distribution of Y by means of the transformation function h is that the likelihood for arbitrary measurements can be written and implemented in an extremely compact form.

For a given transformation function h , the likelihood contribution of a datum $(y, \bar{y}) \in \mathcal{Y}$ is defined in terms of the distribution function (Lindsey 1996), where one can differentiate between four special cases:

$$\mathcal{L}(h \mid Y \in (y, \bar{y}]) \begin{cases} = F_Z(h(\bar{y})) - F_Z(h(y)) & y \in (y, \bar{y}] \quad \text{'interval-censored'} \\ = 1 - F_Z(h(y)) & y \in (y, \infty) \quad \text{'right-censored'} \\ = F_Z(h(\bar{y})) & y \in (-\infty, \bar{y}] \quad \text{'left-censored'} \\ \approx f_Z(h(y))h'(y) & y = (y + \bar{y})/2 \quad \text{'exact continuous'}. \end{cases}$$

For truncated observations in the interval $(y_l, y_r] \subset \mathcal{Y}$, the above likelihood contribution has to be multiplied by the factor $(F_Z(h(y_r)) - F_Z(h(y_l)))^{-1}$ when $y_l < y < \bar{y} \leq y_r$.

For the sake of notional simplicity we only discuss the case of a continuous target Y (see Hothorn *et al.* 2016, for a discussion of the general case), where we always practically observe an imprecise datum $(y, \bar{y}) \subset \mathbb{R}$ and, for short intervals $(y, \bar{y}]$, approximate the exact likelihood by the density $f_Y(y \mid h) = f_Z(h(y))h'(y)$ with $y = (y + \bar{y})/2$ (Lindsey 1999). This approximation only works for relatively precise measurements, i.e., short intervals. If longer intervals are observed, one speaks of “interval censoring” and relies on the definition of the likelihood contribution as the probability of the interval $(y, \bar{y}]$.

To simplify estimation, we parameterise the transformation function $h(y)$ as a linear function of its basis-transformed argument y using a basis function $\mathbf{a} : \mathcal{Y} \rightarrow \mathbb{R}^P$ such that $h(y) = \mathbf{a}(y)^\top \boldsymbol{\vartheta}$, $\boldsymbol{\vartheta} \in \mathbb{R}^P$. In the following, we will write $h = \mathbf{a}^\top \boldsymbol{\vartheta}$. For continuous targets Y the parameterisation $h(y) = \mathbf{a}(y)^\top \boldsymbol{\vartheta}$ needs to be smooth in y , so any polynomial or spline basis is a suitable choice for \mathbf{a} . For the empirical experiments in Section 7 we applied Bernstein polynomials (for an overview see Farouki 2012) of order M ($P = M + 1$) defined on the interval $[\underline{y}, \bar{y}]$ with

$$\begin{aligned} \mathbf{a}_{\text{Bs}, M}(y) &= (M + 1)^{-1} (f_{\text{Be}(1, M+1)}(\tilde{y}), \dots, f_{\text{Be}(m, M-m+1)}(\tilde{y}), \dots, f_{\text{Be}(M+1, 1)}(\tilde{y}))^\top \in \mathbb{R}^{M+1} \\ h(y) &= \mathbf{a}_{\text{Bs}, M}(y)^\top \boldsymbol{\vartheta} = \sum_{m=0}^M \vartheta_m f_{\text{Be}(m+1, M-m+1)}(\tilde{y}) / (M + 1) \\ h'(y) &= \mathbf{a}'_{\text{Bs}, M}(y)^\top \boldsymbol{\vartheta} = \sum_{m=0}^{M-1} (\vartheta_{m+1} - \vartheta_m) f_{\text{Be}(m+1, M-m)}(\tilde{y}) M / ((M + 1)(\bar{y} - \underline{y})) \end{aligned}$$

where $\tilde{y} = (y - \underline{y}) / (\bar{y} - \underline{y}) \in [0, 1]$ and $f_{\text{Be}(m, M)}$ is the density of the Beta distribution with parameters m and M . This choice is computationally attractive because strict monotonicity can be formulated as a set of M linear constraints on the parameters $\vartheta_m < \vartheta_{m+1}$ for all $m = 0, \dots, M$ (Curtis and Ghosh 2011).

The distribution family $\mathbb{P}_{Y,\Theta} = \{F_Z \circ \mathbf{a}^\top \boldsymbol{\vartheta} \mid \boldsymbol{\vartheta} \in \Theta\}$ we define transformation forests upon is called transformation family of distributions with parameter space $\Theta = \{\boldsymbol{\vartheta} \in \mathbb{R}^P \mid \mathbf{a}^\top \boldsymbol{\vartheta} \in \mathcal{H}\}$ and transformation functions $\mathbf{a}^\top \boldsymbol{\vartheta} \in \mathcal{H}$. Examples of the richness of this family are given in Figures 4, 6, 11, 15, or 19 by densities of varying shapes, including various levels of skewness and bimodality. The log-likelihood contribution $\ell_i(\boldsymbol{\vartheta})$ for observation $y_i = (\underline{y}, \bar{y}]_i$ evaluated at some parameter vector $\boldsymbol{\vartheta}$ is now $\ell_i(\boldsymbol{\vartheta}) = \log(\mathcal{L}(\mathbf{a}^\top \boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]_i))$. For the empirical results in Section 7, we always approximated the log-likelihood for an observation $y_i \in \mathbb{R}$ by the log-density using $\ell_i(\boldsymbol{\vartheta}) \approx \log(f_Z(\mathbf{a}(y_i)^\top \boldsymbol{\vartheta})) + \log(\mathbf{a}'(y_i)^\top \boldsymbol{\vartheta})$.

4. Transformation Trees and Forests

Conceptually, the model-based recursive partitioning algorithm (Zeileis *et al.* 2008) for tree induction starts with the maximum likelihood estimator $\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N$ and aims at partitioning the probability model $\mathbb{P}_{Y, \hat{\boldsymbol{\vartheta}}_{\text{ML}}^N}$. Deviations from such a given model that can be explained by parameter instabilities due to one or more of the predictors are investigated based on the score contributions. The novel “transformation trees” suggested here rely on the transformation family $\mathbb{P}_{Y,\Theta} = \{F_Z \circ \mathbf{a}^\top \boldsymbol{\vartheta} \mid \boldsymbol{\vartheta} \in \Theta\}$ whose score contributions

$$\mathbf{s}(\boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]) = \frac{\partial \log(\mathcal{L}(\mathbf{a}^\top \boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]))}{\partial \boldsymbol{\vartheta}}$$

have relatively simple and generic forms. The score contribution of an “exact continuous” observation $y = (\underline{y} + \bar{y})/2$ from an absolutely continuous distribution is approximated by the gradient of the log-density with respect to $\boldsymbol{\vartheta}$

$$\mathbf{s}(\boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]) \approx \mathbf{a}(y) \frac{f'_Z(\mathbf{a}(y)^\top \boldsymbol{\vartheta})}{f_Z(\mathbf{a}(y)^\top \boldsymbol{\vartheta})} + \frac{\mathbf{a}'(y)}{\mathbf{a}'(y)^\top \boldsymbol{\vartheta}}.$$

For an interval-censored observation $(\underline{y}, \bar{y}]$ the score contribution is

$$\mathbf{s}(\boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]) = \frac{f_Z(\mathbf{a}(\bar{y})^\top \boldsymbol{\vartheta})\mathbf{a}(\bar{y}) - f_Z(\mathbf{a}(\underline{y})^\top \boldsymbol{\vartheta})\mathbf{a}(\underline{y})}{F_Z(\mathbf{a}(\bar{y})^\top \boldsymbol{\vartheta}) - F_Z(\mathbf{a}(\underline{y})^\top \boldsymbol{\vartheta})}.$$

For a truncated observation, the score function changes to $\mathbf{s}(\boldsymbol{\vartheta} \mid Y \in (\underline{y}, \bar{y}]) - \mathbf{s}(\boldsymbol{\vartheta} \mid Y \in (y_l, y_r])$.

We start with the model $\mathbb{P}_{Y, \hat{\boldsymbol{\vartheta}}_{\text{ML}}^N}$. The hypothesis of all observations $i = 1, \dots, N$ coming from this model can be written as the independence of the P -dimensional score contributions and all predictors, i.e.,

$$H_0 : \mathbf{s}(\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N \mid Y) \perp \mathbf{X}.$$

This hypothesis can be tested either using asymptotic M-fluctuation tests (Zeileis *et al.* 2008) or permutation tests (Hothorn *et al.* 2006; Zeileis and Hothorn 2013) with appropriate multiplicity adjustment depending on the number of predictors. Rejection of H_0 leads to the implementation of a binary split identified as

$$\hat{\mathcal{B}} = \arg \max_{\mathcal{B} \subset \mathcal{X}} \max_{\boldsymbol{\vartheta}_1, \boldsymbol{\vartheta}_2 \in \Theta} \sum_{i=1}^N I(\mathbf{x}_i \in \mathcal{B}) \ell_i(\boldsymbol{\vartheta}_1) + \sum_{i=1}^N I(\mathbf{x}_i \in \bar{\mathcal{B}}) \ell_i(\boldsymbol{\vartheta}_2).$$

Usually, one restricts partitions $\mathcal{X} = \hat{\mathcal{B}} \cup \bar{\hat{\mathcal{B}}}$ to those obtained from rectangular splits in a single predictor variable, possibly under some sample size constraints. Unbiasedness of the tree with respect to variable selection is a consequence of splitting in the variable of highest association to the scores where association is measured by the marginal multiplicity-adjusted p -value (for details see [Hothorn et al. 2006](#); [Zeileis et al. 2008](#)). The procedure is recursively iterated until H_0 cannot be rejected. The result is a partition of the sample space $\mathcal{X} = \bigcup_{b=1, \dots, B} \mathcal{B}_b$.

Based on the “transformation trees” introduced here, we construct a corresponding random forest-type algorithm as follows. A “transformation forest” is an ensemble of T transformation trees fitted to subsamples of the learning sample and, optionally, a random selection of candidate predictors available for splitting in each node of the tree. The result is a set of T partitions of the predictor sample space. The transformation forest conditional parameter function is defined by its nearest neighbour forest weights (3).

The question arises how the order M of the parameterisation of the transformation function h via Bernstein polynomials affects the conditional distribution functions $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Tree}}^N(\mathbf{x})}$ and $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})}$. On the one hand, the basis $\mathbf{a}_{\text{Bs}, 1}$ with $F_Z = \Phi$ only allows linear transformation functions of a standard normal and thus our models for $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}}$ are restricted to the normal family, however, with potentially both mean and variance depending on \mathbf{x} as the split criterion in transformation trees is sensitive to changes in both mean and variance. This most simple parameterisation leads to transformation trees and forests from which both the conditional mean and the conditional variance can be inferred. Using a higher order M also allows modelling non-normal distributions, in the extreme case with $M = N - 1$ the unconditional distribution function $F_Z(\mathbf{a}_{\text{Bs}, M}(y)^\top \boldsymbol{\theta})$ interpolates the unconditional empirical cumulative distribution function of the target. With $M > 1$, the split criterion introduced in this section is able to detect changes beyond the second moment and, consequently, also higher moments of the conditional distributions $\mathbb{P}_{Y|\mathbf{X}=\mathbf{x}}$ may vary with \mathbf{x} .

5. Transformation Forest Inference

In contrast to all other members from the family of random forest regression models, a transformation forest is a fully-specified parametric model. Thus, we can derive all interesting model inference procedures from well-defined probability models and do not need to fall-back to heuristics. Predictions from transformation models are distributions $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})}$ and we can describe these on the scale of the distribution, quantile, density, hazard, cumulative hazard, expectile and other characterising functions. By far not being comprehensive, we introduce prediction intervals, a unified definition of permutation variable importance, the model-based bootstrap and a test for global independence in this section.

5.1. Prediction Intervals and Outlier Detection

For some yet unobserved target Y under predictors \mathbf{x} , a two-sided $(1 - \alpha)$ prediction interval for $Y \mid \mathbf{X} = \mathbf{x}$ and some $\alpha \in (0, 0.5)$ can be obtained by numerical inversion of the conditional distribution $\mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})}$, for example via

$$\text{PI}_\alpha(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_{\text{Forest}}^N) = \left\{ y \in \mathcal{Y} \mid \alpha/2 < \mathbb{P}_{Y, \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x})}(y) \leq 1 - \alpha/2 \right\}$$

with the property

$$\mathbb{P}_{Y|X=\mathbf{x}}(\text{PI}_\alpha(\mathbf{x} | \boldsymbol{\vartheta})) = 1 - \alpha.$$

The empirical level $\mathbb{P}_{Y|X=\mathbf{x}}(\text{PI}_\alpha(\mathbf{x} | \hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N))$ depends on how well the parameters $\boldsymbol{\vartheta}(\mathbf{x})$ are approximated by the forest estimate $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x})$. If for some observation (y_i, \mathbf{x}_i) the corresponding prediction interval $\text{PI}_\alpha(\mathbf{x}_i | \hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N)$ excludes y_i , one can (at level α) suspect this observed target of being an outlier.

5.2. Permutation Variable Importance

The importance of a variable is defined as the amount of change in the risk function when the association between one predictor variable and the target is artificially broken. Permutation variable importances permute one of the predictors at a time (and thus also break the association to the remaining predictors). The risk function for transformation forests is the (negative) log-likelihood, thus a universally applicable formulation of variable importance for all types of target distributions in transformation forests is

$$\text{VI}(j) = T^{-1} \sum_{t=1}^T \left(\sum_{i=1}^N -\ell_i \left(\hat{\boldsymbol{\vartheta}}_{\text{Tree},t}^N(\mathbf{x}_i) \right) - \sum_{i=1}^N -\ell_i \left(\hat{\boldsymbol{\vartheta}}_{\text{Tree},t}^N(\mathbf{x}_i^{(j)}) \right) \right)$$

where the j -th variable was permuted in $\mathbf{x}_i^{(j)}$ for $i = 1, \dots, N$.

5.3. Model-Based Bootstrap

We use the model-based bootstrap to assess the variability of the estimated forest conditional parameter function $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N$ as follows. First, we fit a transformation forest and sample new target values $\tilde{y}_i \sim \mathbb{P}_{Y, \hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x}_i)}$ for each observation $i = 1, \dots, N$ from this transformation forest. For these $i = 1, \dots, N$ pairs of artificial targets and original predictors $(\tilde{y}_i, \mathbf{x}_i)$, we refit the transformation forest. This procedure of sampling and refitting is repeated $k = 1, \dots, K$ times. The resulting K conditional parameter functions $\hat{\boldsymbol{\vartheta}}_{\text{Forest},k}^N$ are a sample from the distribution of conditional parameter functions assuming the initial $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N$ was the true conditional parameter function. Extreme samples can be identified by looking at the log-likelihood ratio statistic between the estimated and the “true” transformation forest for each of the K bootstrap samples

$$\text{logLR}_k = \sum_{i=1}^N \ell_{i,k} \left(\hat{\boldsymbol{\vartheta}}_{\text{Forest},k}^N(\mathbf{x}_i) \right) - \sum_{i=1}^N \ell_{i,k} \left(\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x}_i) \right).$$

Here, $\ell_{i,k}$ is the log-likelihood contribution by the i -th observation from the k -th bootstrap sample. One can now remove extreme samples (those logLR_k is greater than the empirical $1 - \alpha$ quantile) and study variability of $\hat{\boldsymbol{\vartheta}}_{\text{Forest},k}^N(\mathbf{x})$ for the remaining k and for interesting \mathbf{x} on an appropriate scale.

5.4. Independence Likelihood-Ratio Test

The first question many researchers have is “Is there any signal in my data?”, or, in other words, is the target Y independent of all predictors \mathbf{X} ? Classical tests, such as the F -test in a linear model or multiplicity-adjusted univariate tests, have very low power against

complex alternatives, i.e., in situations where the impact of the predictors is neither linear nor marginally visible. Because transformation forests can potentially detect such structures, we propose a likelihood-ratio test for the null $H_0 : Y \perp \mathbf{X}$. This null hypothesis is identical to $H_0 : \mathbb{P}_Y = \mathbb{P}_{Y|\mathbf{X}=\mathbf{x}} \forall \mathbf{x} \in \mathcal{X}$ and reads $H_0 : \mathbb{P}_{Y,\boldsymbol{\vartheta}} = \mathbb{P}_{Y,\boldsymbol{\vartheta}(\mathbf{x})} \forall \mathbf{x} \in \mathcal{X}$, or even simpler, $H_0 : \boldsymbol{\vartheta}(\mathbf{x}) \equiv \boldsymbol{\vartheta}$ for the class of models we are studying. Under the null hypothesis, the unconditional maximum likelihood estimator $\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N$ would be optimal. It therefore makes sense to compare the log-likelihoods of the unconditional model with the log-likelihood of the transformation forest using the log-likelihood ratio statistic

$$\log\text{LR} = \sum_{i=1}^N \ell_i \left(\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x}_i) \right) - \sum_{i=1}^N \ell_i \left(\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N \right)$$

Under H_0 we expect small differences and under the alternative we expect to see larger log-likelihoods of the transformation forest. The null distribution can be approximated by the model-based bootstrap (Section 5.3). We first estimate the unconditional model $\mathbb{P}_{Y,\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N}$ and, in a second step, draw $k = 1, \dots, K$ samples from this model $\mathbb{P}_{Y,\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N}$ of size N , i.e., we sample under the conditions stated by H_0 (in this sense treating $\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N$ as the “true” parameter). In the k -th sample the predictors are identical to the those in the learning sample and only the target values are replaced. For each of these k samples we refit the transformation forest and obtain $\hat{\boldsymbol{\vartheta}}_{\text{Forest},k}^N(\mathbf{x}_i)$. Based on this model we compute the log-likelihood ratio statistic

$$\log\text{LR}_k = \sum_{i=1}^N \ell_{i,k} \left(\hat{\boldsymbol{\vartheta}}_{\text{Forest},k}^N(\mathbf{x}_i) \right) - \sum_{i=1}^N \ell_{i,k} \left(\hat{\boldsymbol{\vartheta}}_{\text{ML}}^N \right)$$

where $\ell_{i,k}$ is the log-likelihood contribution by the i -th observation from the k -th bootstrap sample. The p -value for H_0 is now $K^{-1} \sum_k I(\log\text{LR}_k > \log\text{LR})$.

6. Theoretical Evaluation

The theoretical properties of random forest-type algorithms are a contemporary research problem and we refer to [Biau and Scornet \(2016\)](#) for an overview. In this section we discuss how these developments relate to the asymptotic behaviour of transformation trees and transformation forests.

For $\boldsymbol{\vartheta}(\mathbf{x}) \equiv \boldsymbol{\vartheta}$ the maximum likelihood estimator ($w_i \equiv 1$) is consistent and asymptotically normal ([Hothorn et al. 2016](#)). In the non-parametric setup, i.e., for arbitrary distributions \mathbb{P}_Y , [Hothorn et al. \(2014\)](#) provide consistency results in the class of conditional transformation models. Based on these results, consistency and normality of the local likelihood estimator for an a priori known partition $\mathcal{X} = \bigcup_{b=1, \dots, B} \mathcal{B}_b$ is guaranteed as long as the sample size tends to infinity in all cells b .

If the partition (transformation trees) or the nearest neighbour weights (transformation forests) are estimated from the data, established theoretical results on random forests ([Lin and Jeon 2006](#); [Meinshausen 2006](#); [Biau et al. 2008](#); [Biau 2012](#); [Scornet et al. 2015](#)) provide a basis for the analysis of transformation forests. [Lin and Jeon \(2006\)](#) first analysed random forests with adaptive nearest neighbours weights, where estimators for the conditional mean

of the form

$$\hat{\mathbb{E}}_N(Y \mid \mathbf{X} = \mathbf{x}) = \frac{\sum_{i=1}^N w_{\text{Forest},i}^N(\mathbf{x}) Y_i}{\sum_{i=1}^N w_{\text{Forest},i}^N(\mathbf{x})}$$

were shown to be consistent in non-adaptive random forests

$$\mathbb{E}_{Y \mid \mathbf{X}=\mathbf{x}} \left(\mathbb{E}(Y \mid \mathbf{X} = \mathbf{x}) - \hat{\mathbb{E}}_N(Y \mid \mathbf{X} = \mathbf{x}) \right)^2 \rightarrow 0$$

as $N \rightarrow \infty$. [Meinshausen \(2006\)](#) showed a Glivenko-Cantelli-type result for conditional distribution functions

$$\hat{\mathbb{P}}_N(Y \leq y \mid \mathbf{X} = \mathbf{x}) = \hat{\mathbb{E}}_N(I(Y \leq y) \mid \mathbf{X} = \mathbf{x}) = \frac{\sum_{i=1}^N w_{\text{randomForest},i}^N(\mathbf{x}) I(Y_i \leq y)}{\sum_{i=1}^N w_{\text{randomForest},i}^N(\mathbf{x})} \quad (4)$$

where the weights are obtained from Breiman's original random forest ([Breiman 2001](#)).

In order to understand the applicability of these results to transformation forests, we define the expected conditional log-likelihood given \mathbf{x} for a fixed set of parameters $\boldsymbol{\vartheta}$ as

$$\ell(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) := \mathbb{E}_{Y \mid \mathbf{X}=\mathbf{x}} \ell(\boldsymbol{\vartheta}, Y),$$

where $\ell(\boldsymbol{\vartheta}, Y_i) = \ell_i(\boldsymbol{\vartheta})$ is the likelihood contribution by some observation Y_i . By definition, the true unknown parameter $\boldsymbol{\vartheta}(\mathbf{x})$ has minimal expected risk and thus maximises the expected log-likelihood, i.e.,

$$\boldsymbol{\vartheta}(\mathbf{x}) = \arg \max_{\boldsymbol{\vartheta} \in \Theta} \ell(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}).$$

Our random forest-type estimator of the expected conditional log-likelihood given \mathbf{x} for a fixed set of parameters $\boldsymbol{\vartheta}$ is now

$$\hat{\ell}_N(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) = \frac{\sum_{i=1}^N w_{\text{Forest},i}^N(\mathbf{x}) \ell(\boldsymbol{\vartheta}, Y_i)}{\sum_{i=1}^N w_{\text{Forest},i}^N(\mathbf{x})}.$$

Under the respective conditions on the distribution of \mathbf{X} and the joint distribution of Y, \mathbf{X} given by [Lin and Jeon \(2006\)](#), [Biau and Devroye \(2010\)](#), or [Biau \(2012\)](#), this estimator is consistent for all $\boldsymbol{\vartheta} \in \Theta$

$$\mathbb{E}_{Y \mid \mathbf{X}=\mathbf{x}} \left(\ell(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) - \hat{\ell}_N(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) \right)^2 \rightarrow 0$$

(the result being derived for non-adaptive random forests). This result gives us consistency of the conditional log-likelihood function

$$\hat{\ell}_N(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) \xrightarrow{\mathbb{P}} \ell(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x}) \quad \forall \boldsymbol{\vartheta} \in \Theta.$$

The forest conditional parameter function $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x})$ is consistent when

$$\mathbb{P}_{\boldsymbol{\vartheta}}(\hat{\ell}_N(\boldsymbol{\vartheta}_1 \mid \mathbf{X} = \mathbf{x}) < \hat{\ell}_N(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x})) \xrightarrow{\mathbb{P}} 1$$

as $N \rightarrow \infty$ for all $\boldsymbol{\vartheta}_1$ in a neighbourhood of $\boldsymbol{\vartheta}$. The result $\hat{\boldsymbol{\vartheta}}_{\text{Forest}}^N(\mathbf{x}) \xrightarrow{\mathbb{P}} \boldsymbol{\vartheta}(\mathbf{x})$ can be shown under the assumptions regarding ℓ given by [Hothorn et al. \(2016\)](#), especially continuity in $\boldsymbol{\vartheta}$.

Because the conditional log-likelihood $\hat{\ell}_N(\boldsymbol{\vartheta} \mid \mathbf{X} = \mathbf{x})$ is a conditional mean-type estimator of a transformed target Y , future theoretical developments in the asymptotic analysis of more realistic random forest-type algorithms based on nearest neighbour weights will directly carry over to transformation forests.

It is worth noting that some authors studied properties of random forests in regression models of the form $Y = r(\mathbf{x}) + \varepsilon$ where the conditional variance $\mathbb{V}(Y \mid \mathbf{X} = \mathbf{x})$ does not depend on \mathbf{x} . This is in line with the ANOVA split criterion implemented in Breiman's original random forest. The split procedure applied in transformation trees is, as will be illustrated in the next section, able to detect changes in higher moments. Thus, transformation forests might be a way to relax the assumption of additivity of signal and noise in the future.

7. Empirical Evaluation

An empirical evaluation of transformation forests and a comparison of this novel member of the random forest family to quantile regression forests (Meinshausen 2006) was performed with the aim to assess the following hypotheses:

- H1:** Transformation trees are able to identify subgroups with differential transformation models, i.e., subgroups defined by a tree whose terminal nodes correspond to transformation models with different parameters and thus different conditional distributions $\mathbb{P}_{Y \mid \mathbf{X}=\mathbf{x}}$.
- H2:** Transformation forests are able to identify differential transformation models where the parameters $\boldsymbol{\vartheta}(\mathbf{x})$ depend on the predictors \mathbf{x} in a more complex way.
- H3:** Quantile regression forests perform as good as transformation forests when higher moments of the conditional distribution $\mathbb{P}_{Y \mid \mathbf{X}=\mathbf{x}}$ are linked to its mean.
- H4:** Transformation forests will outperform quantile regression forests when higher moments of the conditional distribution $\mathbb{P}_{Y \mid \mathbf{X}=\mathbf{x}}$ depend on the predictors, in a way that is not linked to the mean.

In order to test these hypotheses, we performed a data-based simulation experiment with data generating processes derived from models for the Abalone ($n = 4,177$ with 8 predictors), BigMac ($n = 69$ with 9 predictors), Boston Housing ($n = 506$ with 18 predictors), and Ozone ($n = 203$ with 12 predictors) regression problems.

Five models corresponding to five different data generating processes were fitted to each of those four regression problems:

Unconditional: An unconditional transformation model with conditional distribution function $\hat{\mathbb{P}}(Y \leq y \mid \mathbf{X} = \mathbf{x}) = \Phi(\mathbf{a}_{\text{Bs},5}(y)^\top \hat{\boldsymbol{\vartheta}}_{\text{ML}}^N)$, i.e., in fact an unconditional distribution lacking impact of the predictors \mathbf{x} on the target Y ,

Linear: A linear transformation model with conditional distribution function $\hat{\mathbb{P}}(Y \leq y \mid \mathbf{X} = \mathbf{x}) = \Phi(\mathbf{a}_{\text{Bs},5}(y)^\top \hat{\boldsymbol{\vartheta}}_{\text{ML}}^N + \mathbf{x}^\top \hat{\boldsymbol{\beta}}_{\text{ML}}^N)$ allowing for changes in the mean and higher moments as a function of the linear predictor $\mathbf{x}^\top \boldsymbol{\beta}$ only,

QRF: A quantile regression forest (Meinshausen 2016) with conditional distribution function (4),

TRT: A transformation tree with conditional distribution function $\hat{\mathbb{P}}(Y \leq y \mid \mathbf{X} = \mathbf{x}) = \Phi(\mathbf{a}_{\text{Bs},5}(y)^\top \hat{\boldsymbol{\theta}}_{\text{Tree}}^N(\mathbf{x}))$, and

TRF: A transformation forest with conditional distribution function $\hat{\mathbb{P}}(Y \leq y \mid \mathbf{X} = \mathbf{x}) = \Phi(\mathbf{a}_{\text{Bs},5}(y)^\top \hat{\boldsymbol{\theta}}_{\text{Forest}}^N(\mathbf{x}))$, as described in Section 4.

All transformation models parameterise the transformation function $h(y) = \mathbf{a}_{\text{Bs},5}(y)^\top \boldsymbol{\theta}$ in terms of a Bernstein polynomial of order five.

From each of the $5 \times N$ estimates $\hat{\mathbb{P}}(Y \leq y \mid \mathbf{X} = \mathbf{x}_i), i = 1, \dots, N$, we drew 100 bootstrap samples (in a parametric model-based way for the transformation models and conditional non-parametrically for quantile regression forests), leading to 100 independent learning samples from the five different data generating processes for each of the four regression problems. Additional 100 independent validation samples were created in the same way. Because we sampled directly from our initial five models, the distribution of both learning and validation samples were completely known and represent four archetypes of regression relationships, i.e., the absence of a regression function (“Unconditional”), a linear regression function (“Linear”), a tree-structured one (“TRT”), and a potentially complex black-box (“QRF” and “TRF”).

Performance comparison between our five candidate models was implemented as follows. For each of the 100 learning samples, all five models were fitted and the validation sample log-likelihood as well as the validation sample check risk (i.e., the sum of the check function evaluated for each observation) defined by the .005, .025, .05, .5, .95, .975 and .995 quantiles was computed on the 100 validation samples. In total, we obtained $4 \times 5 \times 5 \times 100$ samples from each of the check risk functions and $4 \times 5 \times 4 \times 100$ samples from the validation sample log-likelihoods (because there is no likelihood in quantile regression forests).

[Figure 2 about here.]

7.1. Global Results

We used a normal linear mixed model with random intercepts corresponding to the 100 learning samples in each regression problem to model the mean validation sample likelihood (as it was less skewed than the observed log-likelihood) conditional on the five data generating processes, the five candidate models and the four regression problems. The fixed effects included main effects of data generating process and candidate model as well as their interaction. In addition, the main effect of regression problem and the interaction effect of regression problem and data generating process was modelled as a fixed effect. This model was used to obtain all-pair comparisons of the candidate models performance, separately for each data generating process, averaged over the 100 learning samples in the five regression problems.

Figure 2 displays 95% simultaneous confidence intervals for all candidate model comparisons conditional on the data generating process. For samples drawn from unconditional models, the novel transformation trees and transformation forests performed, on average, as good as the correct unconditional transformation model. Linear transformation models seemed to suffer some overfitting. For samples from the linear transformation model, all models performed better than the unconditional model and the correct linear transformation model outperformed the tree-based procedures. The forest was a little better than the tree in approximating this linear structure. Then a tree structure was used to generate the data,

all conditional models were able to pick-up this signal and both trees and forests performed better than the linear transformation model. The forest seemed to stabilise the trees as it outperformed the correct tree model here. When a forest was used to generate the data, either a quantile regression forest or a transformation forest, all conditional models were again able to detect this signal and the transformation forests performed better than transformation trees and the linear transformation model.

[Figure 3 about here.]

We performed the same linear mixed model analysis separately for each of the check risk functions. Figure 3 shows 95% simultaneous confidence intervals (separately multiplicity-adjusted for each quantile) allowing for a performance comparison of all transformation-based candidate models to quantile regression forests. Confidence intervals larger than zero indicate a superior performance of quantile regression forests while confidence intervals smaller than zero indicate inferior performance of quantile regression forests. For the unconditional data generating process, transformation trees and forests suffered significantly less from overfitting compared to the linear transformation model and quantile regression forests. When the linear transformation model was used to generate learning samples, only the linear model outperformed quantile regression forests while the transformation trees and forests were inferior to quantile regression forests except for the very extreme quantiles (0.005 and 0.995) where the situation was in reverse. Transformation trees and forests outperformed quantile regression forests for extreme quantiles (0.005, 0.025, 0.975, and 0.995) and seemed to perform of par with quantile regression forests for estimating the conditional median. More or less the same picture was observed for data generated from a transformation forest while quantile regression forests outperformed all other methods when the data were generated by this procedure (except for the 0.995 quantile, where transformation forests seemed to be better).

Overall, the results of this simulation study indicate that transformation trees are able to identify subgroups with differential distributions as trees performed better than the unconditional or linear transformation models on tree-based data generating processes (H1). The same is true for transformation forests (H2) which had an edge over transformation trees also when a single tree generated the data. Quantile regression forests outperformed transformation trees and forests on linear structures except for extreme quantiles where transformation forests seem to estimate the tails of the conditional distributions better. This finding is in contrast to hypothesis H3. Transformation forests performed better than quantile regression forests when the data was generated by transformation trees or forests, except for the median where both types of forests seem to be on par (H4).

The two deviations from our initial assumptions, i.e., the superior performance of quantile regression forests on linear problems and the equivalent performance of both forest types with respect to the conditional median, might be explained by a relatively low power for detecting changes in the mean when the tests introduced in Section 4 are applied in transformation trees. In our case, the tests were based on the six score functions of a Bernstein polynomial of order five, whereas quantile regression forests apply a one-dimensional ANOVA split search. The latter procedure seems to be more powerful for detecting changes directly in the mean.

One surprising result of this simulation experiment was the severe overfitting problem observed for quantile regression forests for all quantiles in the unconditional situation. The problem was most pronounced when there was no signal in the data, but also visible in the

tails of the conditional distribution estimated by quantile regression forests in the presence of a signal. For the most extreme 0.005 and 0.995 quantiles, transformation forests performed better than quantile regression forests in all scenarios except when the data were generated by the latter procedure.

We discuss these global findings in more detail based on results obtained for the Ozone regression problem. The results of the remaining three regression problems are given in the Appendix.

7.2. Results Ozone

The density corresponding to the unconditional transformation model for the daily maximum one-hour average ozone reading target is shown in Figure 4. This unconditional distribution was regressed on the predictors \mathbf{x} by the linear transformation model (shown in terms of its conditional quantile functions in Figure 5). The variance of ozone concentration increases with larger values of the linear predictor. The transformation tree with five terminal nodes for the Ozone regression problem is shown in Figure 6. Clearly, both location and scale differ across the five terminal nodes. In addition, the distribution is right-skewed in the two leftmost terminal nodes, left-skewed in the rightmost terminal node and somehow symmetric in the remaining two terminal nodes. The more complex black-box regression relationships represented by quantile regression forests and transformation forests are harder to visualise. We illustrate the regression relationships in these two forests by plots of the conditional quantile function for temperature measured at Sandford, CA, while all other predictors were chosen constant at their median or mode in Figure 7. The overall pattern in these conditional quantile functions fitted by the two forest procedures looks quite similar. The most important difference between the two plots is the variance heterogeneity that seems present in the transformation forest while quantile regression forest predicted a more or less constant variance across the temperature range. The permutation variable importance of the original implementation of random forests (Breiman 2001; Breiman *et al.* 2015, as **quantregForest** does not implement a variable importance measure) ranks the temperature measured at El Monte, CA, the temperature measured at Sandburg, CA, and the inversion base temperature at LAX on the top-three positions. The likelihood-based permutation variable importance obtained from transformation forests (Section 5.2) agrees, except that it regards the temperature at Sandburg more important than the temperature at El Monte.

[Figure 4 about here.]

[Figure 5 about here.]

[Figure 6 about here.]

[Figure 7 about here.]

For each of the 100 bootstrap learning samples obtained from each of the five models discussed above, the validation sample log-likelihood (computed on the validation samples) is depicted for each of the candidate models used to fit the bootstrap samples in Figure 8. Since it was not possible to evaluate the likelihood for quantile regression forests, this method is

missing on the x-axis and each panel corresponds to one data generating process. Overall, the corresponding correct candidate model performed best, i.e., lead to the minimum negative log-likelihood, except for single tree-based datasets where the transformation forests performed better than the single transformation tree. This pattern was also observable in the remaining panels, except the panel corresponding to the unconditional model. Here, single transformation trees performed as good as the correct unconditional model, due to the in-built test-based pre-pruning. Under the null of no association between predictors and the target, transformation trees will incorrectly split the root node in not more than 5% of the learning samples (Hothorn *et al.* 2006). Therefore, for the majority of the learning samples, the transformation model in the root node is equivalent to the unconditional transformation model. When the predictors are informative, transformation trees seemed to be less stable than the remaining candidate models as the variance of the validation log-likelihood was largest for transformation trees. Transformation forests suffer overfitting, but to a much lesser extent than linear transformation models.

[Figure 8 about here.]

The results regarding the validation sample check risk support the general finding that the correct models always performed superior or at least on par to its competitors. Quantile regression forests seemed to be superior to transformation forests on linear problems (but not for extreme quantiles) when the linear predictor determines both the mean and higher moments. When changes in higher moments were not linked to changes in the mean, i.e., in the samples drawn from the transformation tree and the transformation forest, transformation forests performed better than quantile regression forests. Also, more extreme quantiles seemed to be better estimated by transformation forests. In the unconditional situation, quantile regression forests suffered from overfitting. In summary, the more formal global findings presented in Subsection 7.1 are in line with this visual inspection of the results obtained from the Ozone regression problem.

[Figure 9 about here.]

8. Discussion

Transformation forests, as well as the underlying transformation trees, can be understood as adaptive local likelihood estimators in the rather general parametric transformation family of distributions. Owing to possible interactions and non-linear effects in a “black-box” conditional parameter function $\vartheta(\mathbf{x})$, the resulting conditional distributions of the target may depend on the predictors in a very general way. The ability to model the impact of some predictors on the whole conditional distribution simultaneously, including its mean and higher moments, is a unique feature of this novel member of the random forest family.

Understanding trees and forests as methods for the estimation of fully parametric conditional distributions allows a much broader set of model inference procedures to be directly applied. The empirical results presented in this paper were based on simulation models obtained from “real” problems (as far as one can understand benchmarking problems as being real). The ability to sample from transformation models helped us to construct data-based instead

of artificial simulation models. Investigations into other inference procedures sketched in Section 5 are currently under way. The likelihood approach taken here also directly allows the procedures to be applied to randomly censored or truncated observations. For example, the rugs in Figure 4 suggest that the ozone readings are in fact interval-censored and the application of the exact instead of the approximate likelihood might be more appropriate. The same applies to right-censored median housing values larger or equal to 50 in the Boston Housing regression problem. A comparison with established trees and forests for the analysis of survival times is an ongoing research project.

The algorithmic internals of transformation trees are rooted in conditional inference trees (Hothorn *et al.* 2006) and model-based recursive partitioning (Zeileis *et al.* 2008) and inherit the unbiased variable selection property from these ancestors. Transformation forests also allow for unbiased variable importances (Strobl *et al.* 2007), including the internal handling of missing predictor variables (Hapfelmeier *et al.* 2014). An open-source implementation of transformation trees and transformation forests based on the **partykit** add-on package (Hothorn and Zeileis 2015) to the R system for statistical computing is available as add-on package **trtf** (Hothorn 2016c).

Within the theory of adaptive local likelihood estimation, alternative choices for weights and parametric models are possible. In the context of personalised medicine or personalised marketing, one is interested in the dependency of some treatment effect β on predictors \mathbf{x} . Random forest-type algorithms are a promising tool for modelling complex effects of predictors on such a treatment parameter (Foster *et al.* 2011; Seibold *et al.* 2016b; Wager and Athey 2015; Seibold *et al.* 2016a). In the framework presented here, implementation of such a strategy only requires the specification of a distribution $\mathbb{P}(Y \leq y \mid \text{treated}) = F_Z(\mathbf{a}(y)^\top \boldsymbol{\vartheta} + \beta)$ for treated and $\mathbb{P}(Y \leq y \mid \text{untreated}) = F_Z(\mathbf{a}(y)^\top \boldsymbol{\vartheta})$ for untreated observations. The model, and therefore also the treatment effect β , can then be partitioned or aggregated by transformation trees and transformation forests leading to a random forest estimate $\hat{\beta}_{\text{Forest}}^N(\mathbf{x})$ of the conditional treatment effect $\beta(\mathbf{x})$. If the parameter estimates for such, or a more complex, model are already given, one could use transformation forests to estimate a deviation from this initial model. An already existing transformation function can be used as an offset in the likelihood ℓ_i , such that the forest conditional parameter function excludes these existing effects.

A more general understanding of weights could be derived from the notion of applying a distance measure d to two distributions $\hat{\mathbb{P}}_{Y|\mathbf{X}=\mathbf{x}} = \hat{\mathbb{P}}_{Y, \hat{\boldsymbol{\vartheta}}_t(\mathbf{x})}$ and $\hat{\mathbb{P}}_{Y|\mathbf{X}=\mathbf{x}_i} = \hat{\mathbb{P}}_{Y, \hat{\boldsymbol{\vartheta}}_t(\mathbf{x}_i)}$ obtained from the t -th tree. Based on this distance, an alternative weight could be defined by

$$w_{\text{Forest},i}^N(\mathbf{x}) = \sum_{t=1}^T (1 - d(\hat{\boldsymbol{\vartheta}}_{\text{Tree},t}^N(\mathbf{x}), \hat{\boldsymbol{\vartheta}}_{\text{Tree},t}^N(\mathbf{x}_i)))$$

for example using the Kullback-Leibler divergence for continuous distributions

$$d_{\text{KL}}(\boldsymbol{\vartheta}_1, \boldsymbol{\vartheta}_2) = \int f_Y(y \mid \boldsymbol{\vartheta}_1) \log \left(\frac{f_Y(y \mid \boldsymbol{\vartheta}_1)}{f_Y(y \mid \boldsymbol{\vartheta}_2)} \right) dy$$

(after standardisation to the unit interval). This weight takes the conditional distribution in two terminal nodes of a tree into account, rather than just treating them as “somehow different” in the way of nearest neighbour weights.

The empirical results suggest that the primary domain of transformation forests should be problems where higher moments are not directly linked to the mean or when one is interested

in more extreme quantiles of the conditional distributions. The only practical problem arises with computational speed, as the computation of the maximum likelihood estimators and the score functions in each node of the underlying trees and the computation of the forest conditional parameter function (3) are more time consuming than the tree-growing and aggregation in quantile regression forests. To speed up computations, one can either reduce the number of parameters P in $\boldsymbol{\vartheta}$ or the number of weights and therefore the number of summands N when computing the conditional forest log-likelihood in (3). The first strategy means that one has to fall-back to simpler distributions $\mathbb{P}_{Y,\boldsymbol{\vartheta}}$ where the maximum likelihood estimator can be computed analytically. Thus, saving computation time would be considered more important than generality of the corresponding conditional distributions. The second option is fast for general conditional distributions owing to the ability of transformation models to deal with interval-censored targets. If one bins the targets y_i into $J + 1$ bins at breaks $-\infty < y_{(1)} < \dots < y_{(J)} < \infty$, a model $\mathbb{P}_{Y,\boldsymbol{\vartheta}}$ of higher complexity can be fitted by maximising the weighted log-likelihood for interval-censored observations $\log(F_Z(h(y_{(j)})) - F_Z(h(y_{(j-1)})))$ when $y_{(j-1)} < y_i \leq y_{(j)}$. In combination with binned predictor variables, considerable improvements with respect to computing time and memory consumption are possible. An experimental implementation of this approach is currently under evaluation.

References

- Biau G (2012). “Analysis of a Random Forests Model.” *Journal of Machine Learning Research*, **13**(1), 1063–1095.
- Biau G, Devroye L (2010). “On the Layered Nearest Neighbour Estimate, the Bagged Nearest Neighbour Estimate and the Random Forest Method in Regression and Classification.” *Journal of Multivariate Analysis*, **101**(10), 2499–2518. doi:10.1016/j.jmva.2010.06.019.
- Biau G, Devroye L, Lugosi G (2008). “Consistency of Random Forests and Other Averaging Classifiers.” *Journal of Machine Learning Research*, **9**, 2015–2033.
- Biau G, Scornet E (2016). “A Random Forest Guided Tour.” *Test*, **25**(2), 197–227. doi:10.1007/s11749-016-0481-7.
- Breiman L (2001). “Random Forests.” *Machine Learning*, **45**(1), 5–32. doi:10.1023/A:1010933404324.
- Breiman L, Cutler A, Liaw A, Wiener M (2015). *Breiman and Cutler’s Random Forests for Classification and Regression*. R package vignette version 4.6-12, URL <https://CRAN.R-project.org/package=randomForest>.
- Brillinger DR (1977). “Discussion of Stone (1977).” *The Annals of Statistics*, **5**(4), 622–623.
- Curtis SM, Ghosh SK (2011). “A Variable Selection Approach to Monotonic Regression with Bernstein Polynomials.” *Journal of Applied Statistics*, **38**(5), 961–976. doi:10.1080/02664761003692423.
- Farouki RT (2012). “The Bernstein Polynomial Basis: A Centennial Retrospective.” *Computer Aided Geometric Design*, **29**(6), 379–419. doi:10.1016/j.cagd.2012.03.001.

- Foster JC, Taylor JM, Ruberg SJ (2011). “Subgroup Identification From Randomized Clinical Trial Data.” *Statistics in Medicine*, **30**(24), 2867–2880. doi:[10.1002/sim.4322](https://doi.org/10.1002/sim.4322).
- Hapfelmeier A, Hothorn T, Ulm K, Strobl C (2014). “A New Variable Importance Measure for Random Forests with Missing Data.” *Statistics and Computing*, **24**(1), 21–34. doi:[10.1007/s11222-012-9349-1](https://doi.org/10.1007/s11222-012-9349-1).
- Hothorn T (2016a). *mlt: Most Likely Transformations*. R package version 0.1-2, URL <http://CRAN.R-project.org/package=mlt>.
- Hothorn T (2016b). *Most Likely Transformations: The mlt Package*. R package vignette version 0.1-4, URL <https://CRAN.R-project.org/package=mlt.docreg>.
- Hothorn T (2016c). *Transformation Trees and Forests*. R package version 0.1-0, URL <https://R-forge.R-project.org/projects/ctm>.
- Hothorn T, Hornik K, Zeileis A (2006). “Unbiased Recursive Partitioning: A Conditional Inference Framework.” *Journal of Computational and Graphical Statistics*, **15**(3), 651–674. doi:[10.1198/106186006X133933](https://doi.org/10.1198/106186006X133933).
- Hothorn T, Kneib T, Bühlmann P (2014). “Conditional Transformation Models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **76**(1), 3–27. doi:[10.1111/rssb.12017](https://doi.org/10.1111/rssb.12017).
- Hothorn T, Lausen B, Benner A, Radespiel-Tröger M (2004). “Bagging Survival Trees.” *Statistics in Medicine*, **23**(1), 77–91. doi:[10.1002/sim.1593](https://doi.org/10.1002/sim.1593).
- Hothorn T, Möst L, Bühlmann P (2016). “Most Likely Transformations.” *Technical report*, arXiv 1508.06749, v2. URL <https://arxiv.org/abs/1508.06749>.
- Hothorn T, Zeileis A (2015). “partykit: A Modular Toolkit for Recursive Partytioning in R.” *Journal of Machine Learning Research*, **16**, 3905–3909. URL <http://jmlr.org/papers/v16/hothorn15a.html>.
- Ishwaran H, Kogalur UB, Blackstone EH, Lauer MS (2008). “Random Survival Forests.” *The Annals of Applied Statistics*, **2**(3), 841–860. doi:[10.1214/08-aos169](https://doi.org/10.1214/08-aos169).
- Lin Y, Jeon Y (2006). “Random Forests and Adaptive Nearest Neighbors.” *Journal of the American Statistical Association*, **101**(474), 578–590. doi:[10.1198/016214505000001230](https://doi.org/10.1198/016214505000001230).
- Lindsey JK (1996). *Parametric Statistical Inference*. Clarendon Press, Oxford.
- Lindsey JK (1999). “Some Statistical Heresies.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, **48**(1), 1–40.
- Loader C (1999). *Local Regression and Likelihood*. Springer-Verlag, New York.
- Meinshausen N (2006). “Quantile Regression Forests.” *Journal of Machine Learning Research*, **7**, 983–999.
- Meinshausen N (2016). *quantregForest: Quantile Regression Forests*. R package version 1.3-5, URL <https://CRAN.R-project.org/package=quantregForest>.

- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Scornet E, Biau G, Vert JP (2015). “Consistency of Random Forests.” *The Annals of Statistics*, **43**(4), 1716–1741. doi:[10.1214/15-AOS1321](https://doi.org/10.1214/15-AOS1321).
- Seibold H, Zeileis A, Hothorn T (2016a). “Individual Treatment Effect Prediction for ALS Patients.” *Technical report*, arXiv 1604.08720. URL <http://arxiv.org/abs/1604.08720>.
- Seibold H, Zeileis A, Hothorn T (2016b). “Model-Based Recursive Partitioning for Subgroup Analyses.” *International Journal of Biostatistics*, **12**(1), 45–63. doi:[10.1515/ijb-2015-0032](https://doi.org/10.1515/ijb-2015-0032).
- Stone CJ (1977). “Consistent Nonparametric Regression (with Discussion).” *The Annals of Statistics*, **5**(4), 595–645. doi:[10.1214/aos/1176343886](https://doi.org/10.1214/aos/1176343886).
- Strobl C, Boulesteix AL, Zeileis A, Hothorn T (2007). “Bias in Random Forest Variable Importance Measures: Illustrations, Sources and a Solution.” *BMC Bioinformatics*, **8**, 25. doi:[10.1186/1471-2105-8-25](https://doi.org/10.1186/1471-2105-8-25).
- Tibshirani R, Hastie T (1987). “Local Likelihood Estimation.” *Journal of the American Statistical Association*, **82**(398), 559–567. doi:[10.1080/01621459.1987.10478466](https://doi.org/10.1080/01621459.1987.10478466).
- Wager S, Athey S (2015). “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Technical report*, arXiv 1510.04342, v2. URL <http://arxiv.org/abs/1510.04342>.
- Zeileis A, Hothorn T (2013). “A Toolbox of Permutation Tests for Structural Change.” *Statistical Papers*, **54**(4), 931–954. doi:[10.1007/s00362-013-0503-4](https://doi.org/10.1007/s00362-013-0503-4).
- Zeileis A, Hothorn T, Hornik K (2008). “Model-Based Recursive Partitioning.” *Journal of Computational and Graphical Statistics*, **17**(2), 492–514. doi:[10.1198/106186008X319331](https://doi.org/10.1198/106186008X319331).

Appendix

Computational Details

A reference implementation of `transformation trees` and `transformation forests` is available in the `trtf` package (Hothorn 2016c). Transformation models were fitted using the `mlt` package (Hothorn 2016a,b) and quantile regression forests were computed by the `quantregForest` package (Meinshausen 2016). All computations were performed using R version 3.3.2 (R Core Team 2016).

For the empirical evaluation in Section 7, all transformation models were based on transformation functions parameterised in terms of Bernstein polynomials of order five. Log-likelihoods were optimised under monotonicity constraints using a combination of augmented Lagrangian minimisation and spectral projected gradients. Transformation trees stopped when the minimum Bonferroni-adjusted p -value was larger than 0.05. No such pre-pruning was applied in transformation forests. The minimum number of observations in each terminal node was 100 (Abalone), 40 (Boston Housing), and 20 (BigMac and Ozone) and nodes

were possibly split when the number of observations was twice as large. Potential censoring in the targets was ignored and the log-likelihood was approximated by the log-density $\ell_i(\boldsymbol{\vartheta}) \approx \log(f_Z(\mathbf{a}(y_i)^\top \boldsymbol{\vartheta})) + \log(\mathbf{a}'(y_i)^\top \boldsymbol{\vartheta})$.

The data analyses presented in this paper can be reproduced from the R code

```
> install.packages(c("mlt", "libcoin", "inum", "partykit", "alr3",
+                   "AppliedPredictiveModeling", "mlbench",
+                   "quantregForest"))
> ### install trtf 0.1-0 and partykit 2.0-0 locally
> library("trtf")
> demo("applications")
```

Development versions of packages **partykit** and **trtf** are temporarily available from

http://user.math.uzh.ch/hothorn/docs/trtf_packages.zip

Additional Results Empirical Evaluation

The linear transformation models for the Abalone, BigMac, and Boston Housing regression problems are given in Figures 10, 14, and 18. For all three regression problems, the conditional variance increased with the linear predictor. The three transformation trees are shown in Figures 11, 15, and 19 with location, scale, and skewness differing between the terminal nodes. The size of the trees differ due to the different sample sizes in these three regression problems.

The validation log-likelihoods for the 100 learning samples in each regression problem are given in Figures 12, 16, and 20. The overall pattern coincides with the global results and the detailed results obtained for the Ozone regression problem. The same applies to the validation check risks in Figures 13, 17, and 21.

[Figure 10 about here.]

[Figure 11 about here.]

[Figure 12 about here.]

[Figure 13 about here.]

[Figure 14 about here.]

[Figure 15 about here.]

[Figure 16 about here.]

[Figure 17 about here.]

[Figure 18 about here.]

[Figure 19 about here.]

[Figure 20 about here.]

[Figure 21 about here.]

Affiliation:

Torsten Hothorn
Institut für Epidemiologie, Biostatistik und Prävention
Universität Zürich
Hirschengraben 84
CH-8001 Zürich, Switzerland
Torsten.Hothorn@uzh.ch

Achim Zeileis
Department of Statistics
Faculty of Economics and Statistics
Universität Innsbruck
Universitätsstraße 15
A-6020 Innsbruck, Austria




Figure 1: Empirical Illustration. Conditional on a uniform predictor x , the distribution $Y \sim \mathcal{N}(0, (1 + 2I(x > .5))^2)$ features a variance breakpoint at 0.5 for 10,000 observations. The red solid lines depict the true conditional 20% and 80% quantiles and the blue dashed lines their estimates obtained from a quantile regression forest (QRF, dashed) and the newly proposed transformation forests (TRF, dotted).




Figure 2: Empirical Evaluation. Comparison of the validation sample likelihood between unconditional and linear transformation models as well as transformation trees (TRT) and forests (TRF) for five different data generating processes, one in each panel. QRF indicates data drawn from a quantile regression forest. 95% simultaneous confidence intervals for the mean negative likelihood difference were obtained from a normal linear mixed model. The horizontal red line indicates equivalent performance. Confidence intervals smaller than zero indicate superior performance of model A in an A - B comparison.

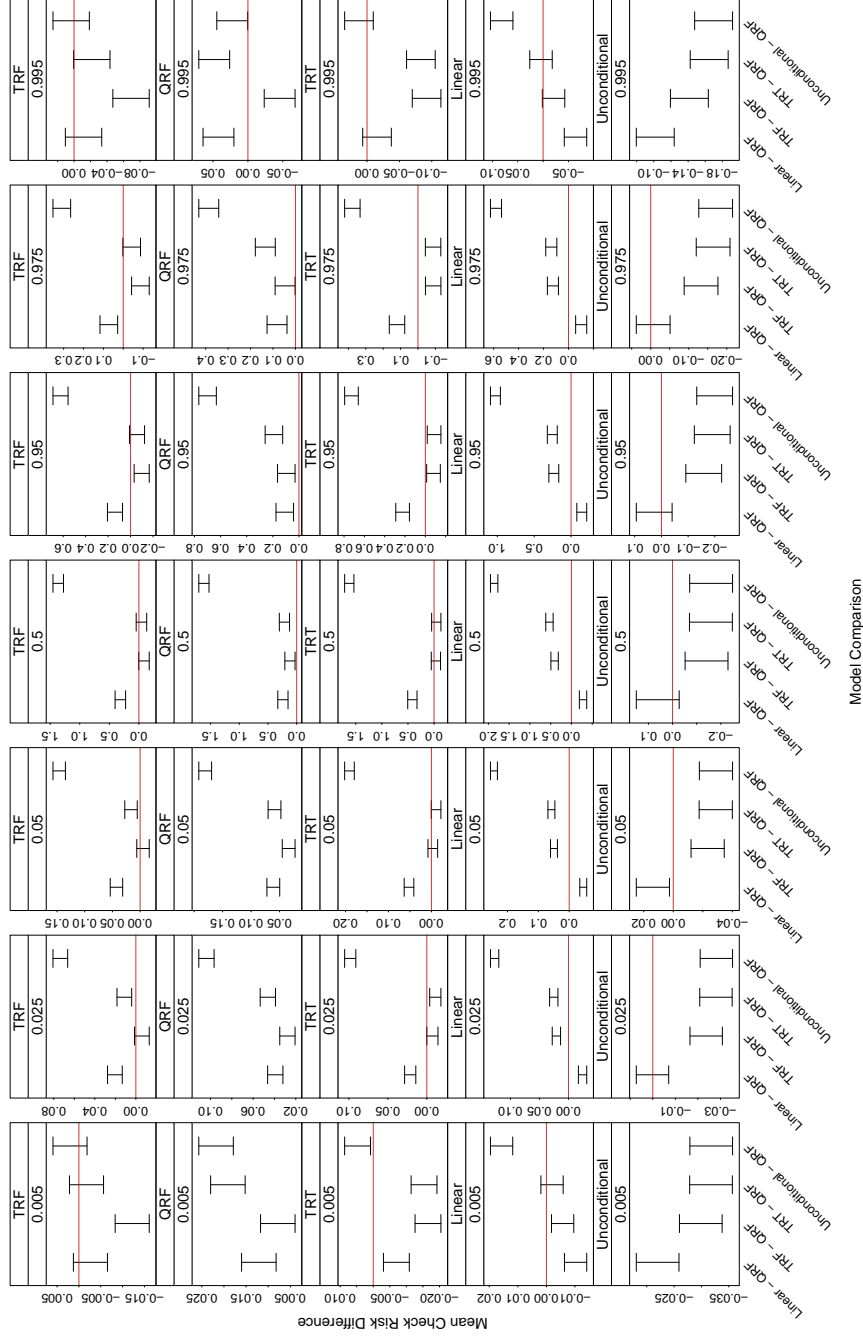


Figure 3: Empirical Evaluation. Comparison of the validation sample check risk for various quantiles (.005, .025, .05, .5, .95, .975, .995) corresponding to columns for quantile regression forests (QRF), unconditional and linear transformation models as well as transformation trees (TRT) and forests (TRF) for five different data generating processes corresponding to rows. 95% simultaneous confidence intervals for the risk difference were obtained from a normal linear mixed model. The horizontal red line indicates equivalent performance. Confidence intervals smaller than zero indicate performance better than quantile regression forests.




Figure 4: Ozone Regression Problem. Unconditional transformation model depicted in terms of the unconditional density estimate.




Figure 5: Ozone Regression Problem. Visualisation of the linear transformation model fitted to the Ozone problem. The value of the linear predictor $\mathbf{x}^\top \hat{\boldsymbol{\beta}}_{\text{ML}}^N$ is plotted against the observations. The lines present the conditional quantile functions for the .1, .2, . . . , .9 quantiles. Larger values of the linear predictor are linked to larger variability.

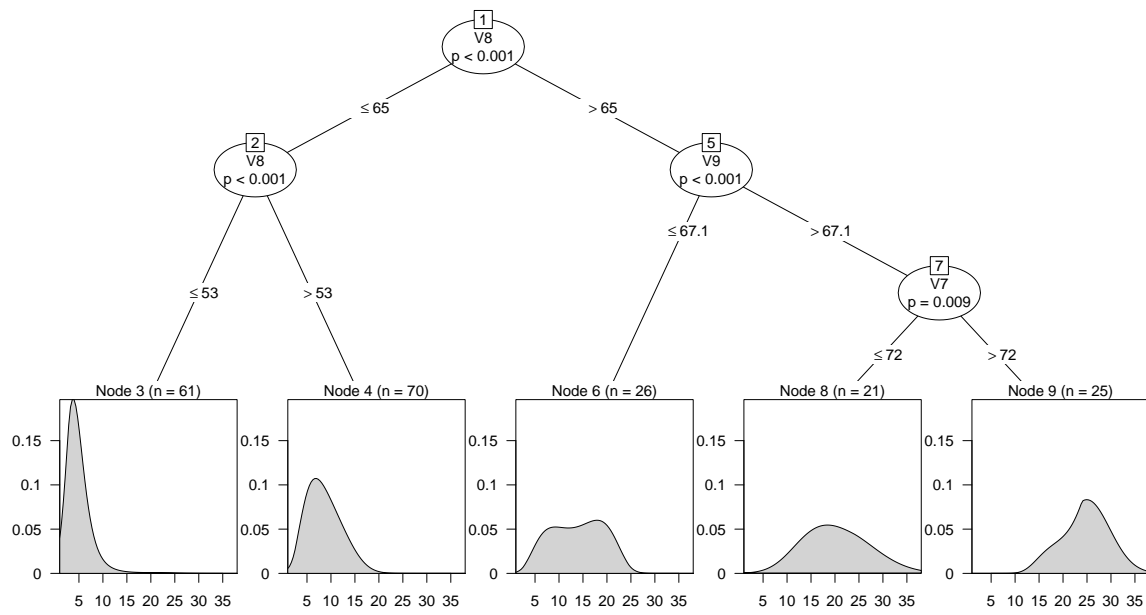


Figure 6: Ozone Regression Problem. Visualisation of transformation tree with conditional densities at the terminal nodes. Location, variance and skewness varies between the five terminal nodes.




Figure 7: Ozone Regression Problem. Conditional quantiles (solid 0.2, dashed median, dotted 0.8) as a function of temperature at Sandburg, CA, for quantile regression forests (QRF) and transformation forests (TRF). All other predictors were chosen constant at their median (numeric measurements) or mode (categorical measurements). The conditional variance seems to be rather stable in quantile regression forests whereas variance heterogeneity can be observed for the transformation forests.




Figure 8: Empirical Evaluation Ozone. Average negative log-likelihood for four transformation models computed on 100 validation samples of size N from five data generating processes corresponding to the five panels. The data generating processes were estimated on the original Ozone data. QRF indicates data drawn from a quantile regression forest, TRT data from a transformation tree and TRF from a transformation forest.




Figure 9: Empirical Evaluation Ozone. Average check risk for varying quantiles (in columns) for unconditional and linear transformation models as well as transformation trees (TRT) and forests (TRF) and quantile regression forests (QRF) computed on 100 validation samples of size N from five data generating processes corresponding to rows. The data generating processes were estimated on the original Ozone data.




Figure 10: Abalone Regression Problem. Visualisation of linear transformation model. The value of the linear predictor $\mathbf{x}^\top \hat{\boldsymbol{\beta}}$ is plotted against the observations. The lines present the conditional quantile functions for the .1, .2, ..., .9 quantiles.

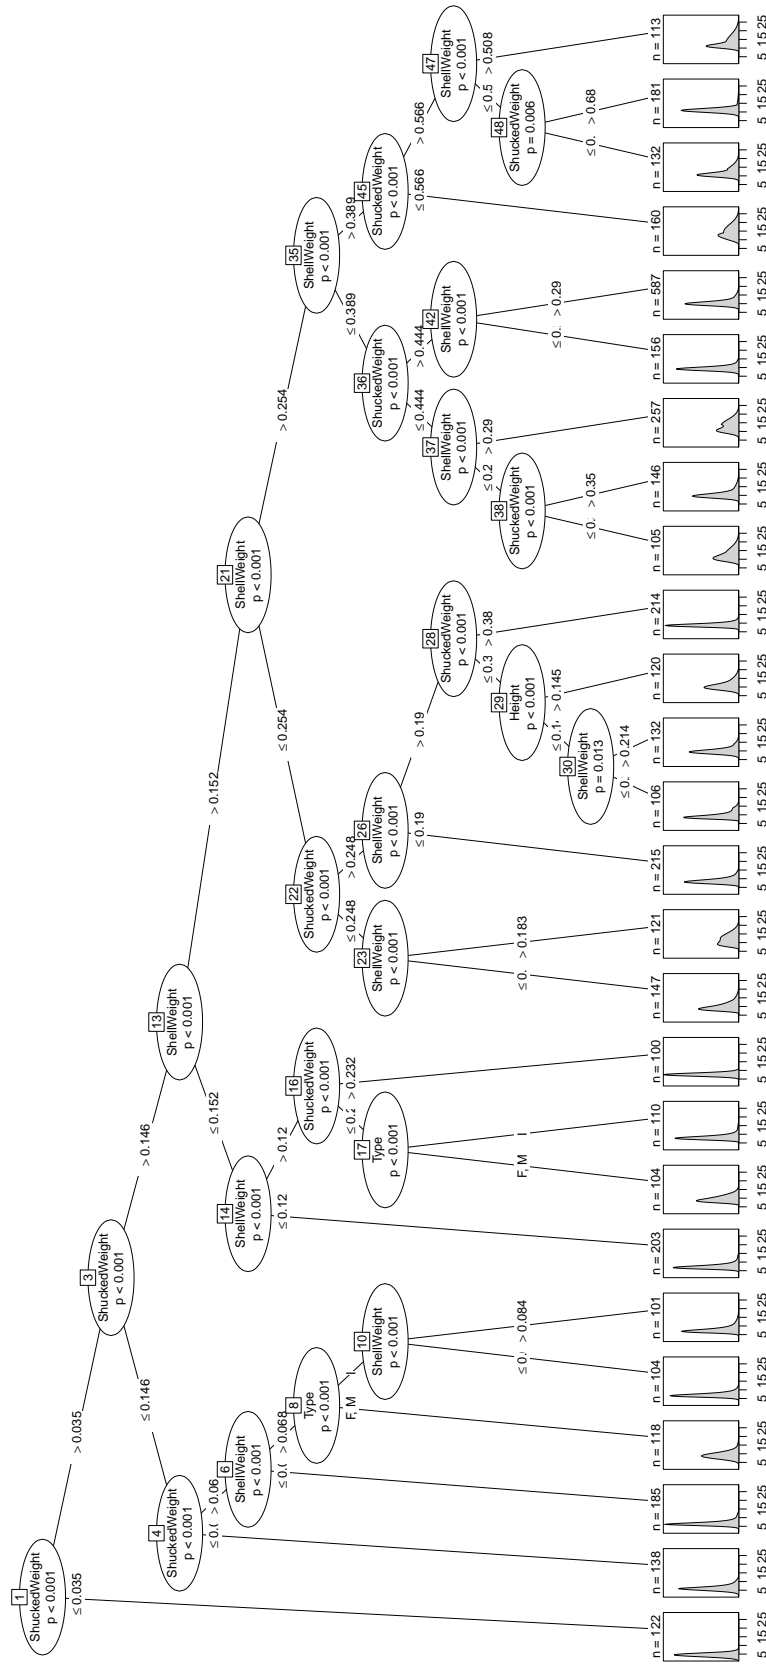


Figure 11: Abalone Regression Problem. Visualisation of transformation tree with conditional densities in the terminal nodes.

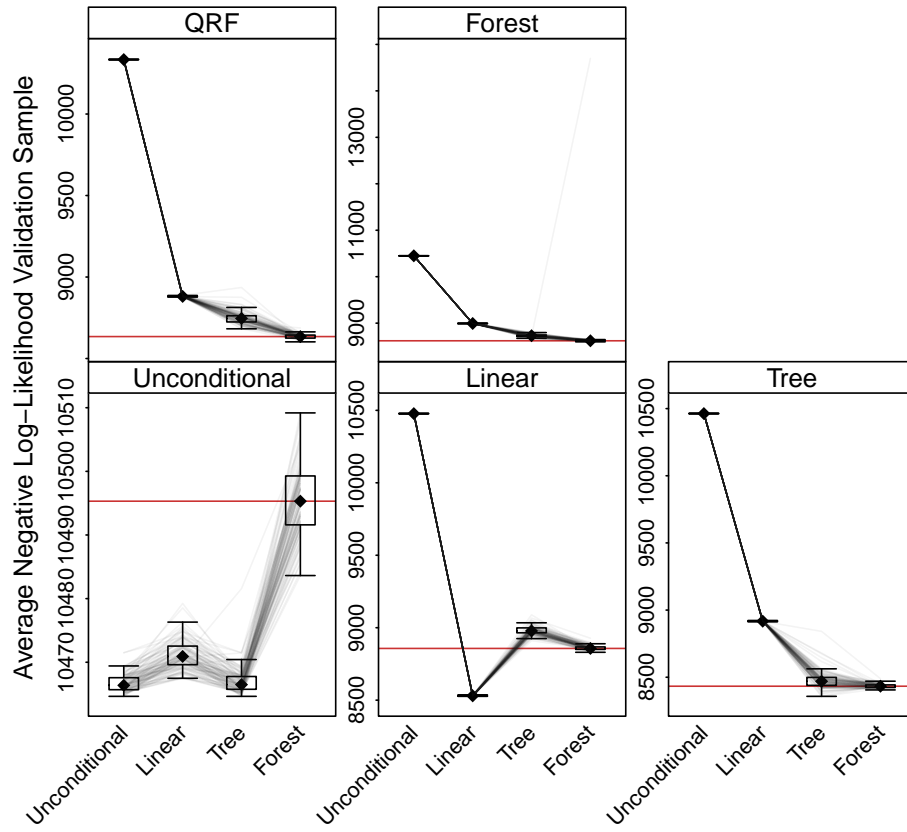


Figure 12: Abalone Regression Problem. Average negative log-likelihood for four transformation models computed on 100 validation samples of size N from five data generating processes corresponding to the five panels. The data generating processes were estimated on the original Abalone data.

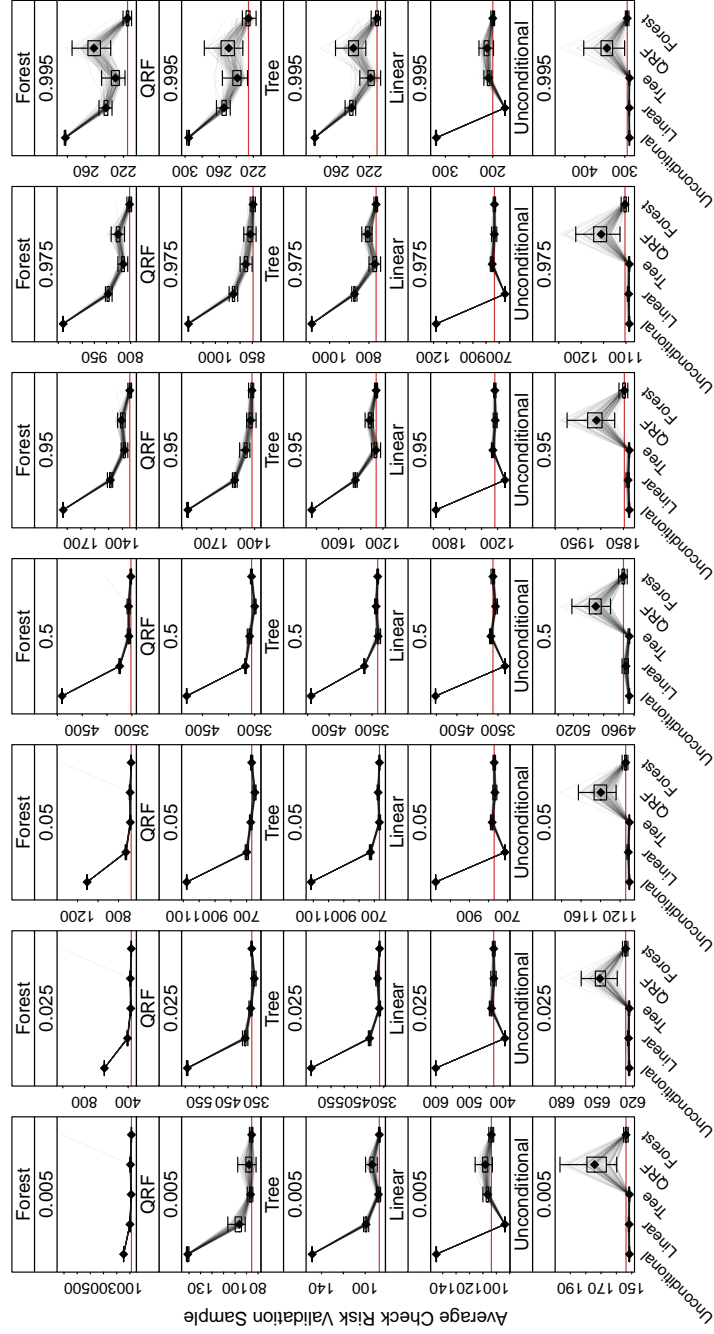


Figure 13: Abalone Regression Problem. Average check risk for varying quantiles (in columns) for four transformation models and quantile regression forests (QRF) computed on 100 validation samples of size N from five data generating processes corresponding to rows. The data generating processes were estimated on the original Abalone data.




Figure 14: BigMac Regression Problem. Visualisation of linear transformation model. The value of the linear predictor $\mathbf{x}^\top \hat{\boldsymbol{\beta}}$ is plotted against the observations. The lines present the conditional quantile functions for the .1, .2, . . . , .9 quantiles.

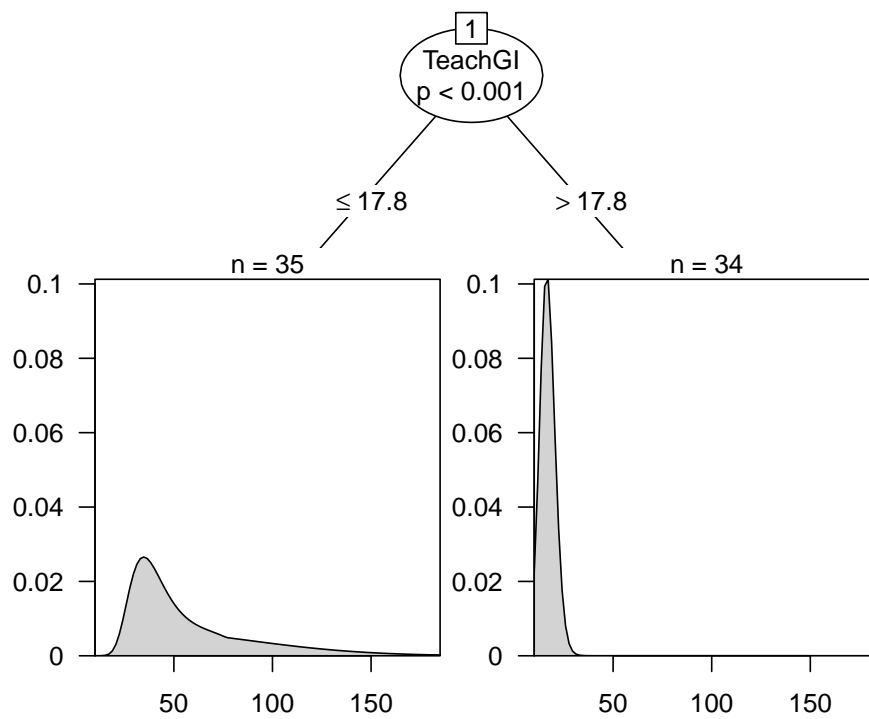


Figure 15: BigMac Regression Problem. Visualisation of transformation tree with conditional densities in the terminal nodes.

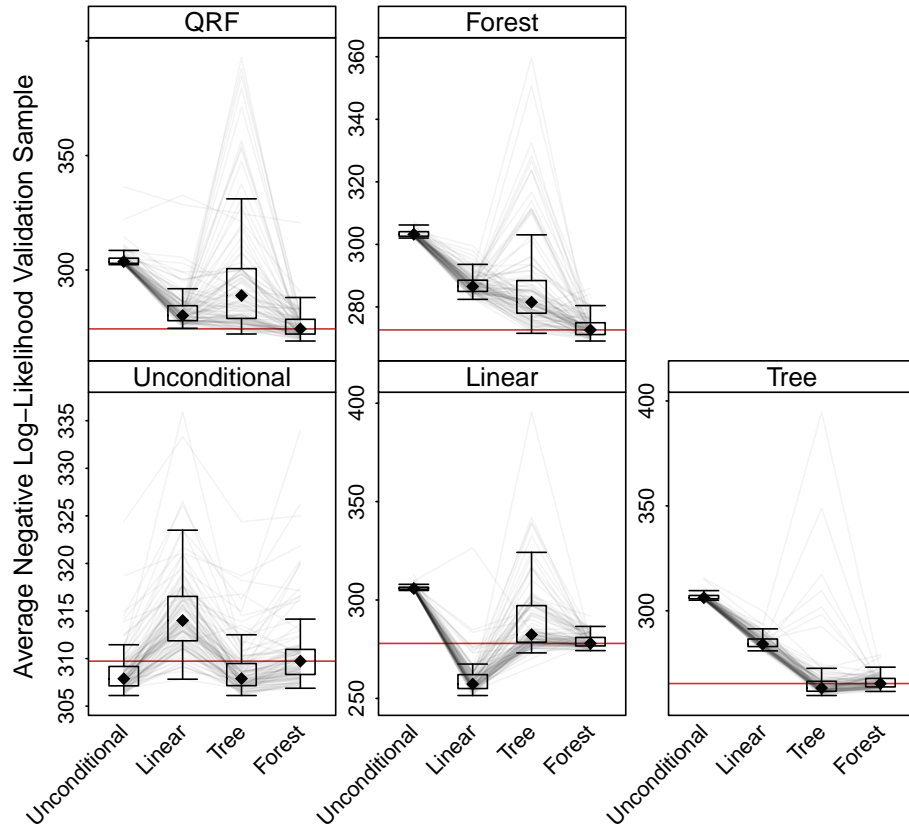


Figure 16: BigMac Regression Problem. Average negative log-likelihood for four transformation models computed on 100 validation samples of size N from five data generating processes corresponding to the five panels. The data generating processes were estimated on the original BigMac data. For better display, 19 extreme outliers were removed from the plot.

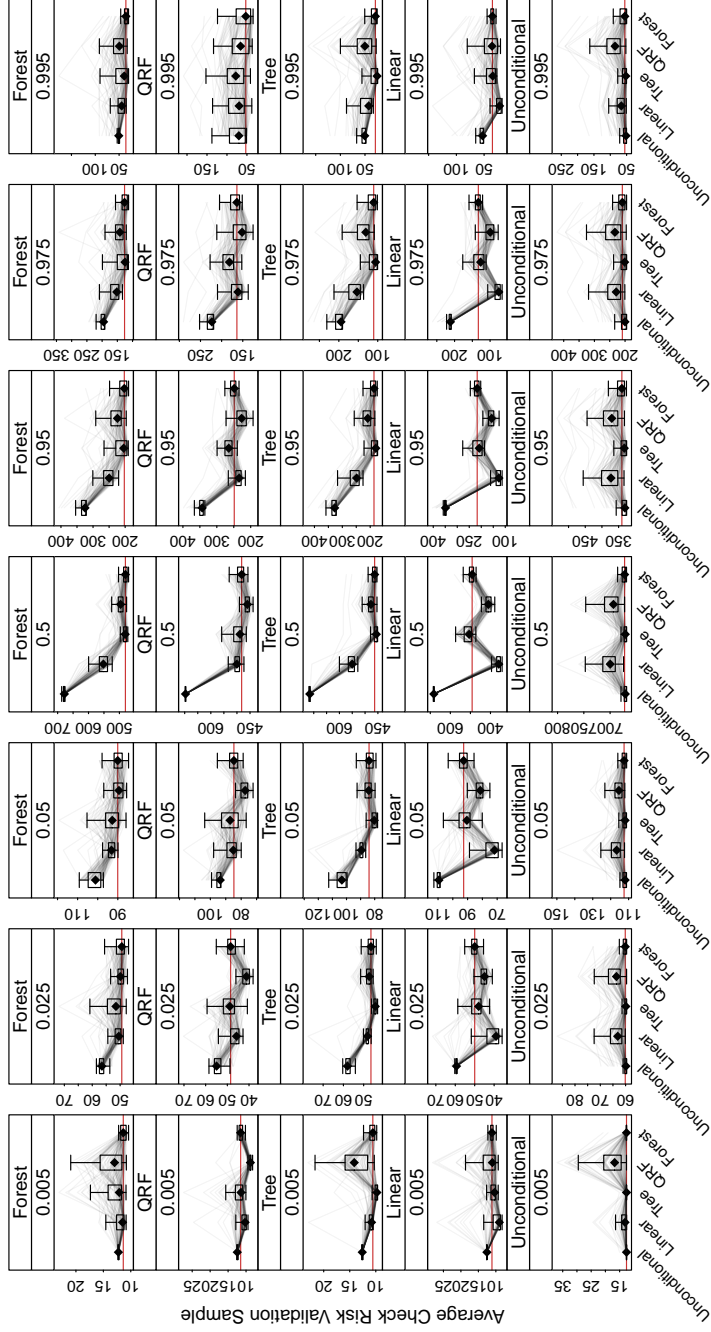


Figure 17: BigMac Regression Problem. Average check risk for varying quantiles (in columns) for four transformation models and quantile regression forests (QRF) computed on 100 validation samples of size N from five data generating processes corresponding to rows. The data generating processes were estimated on the original BigMac data.




Figure 18: Boston Housing Regression Problem. Visualisation of linear transformation model. The value of the linear predictor $\mathbf{x}^\top \hat{\boldsymbol{\beta}}$ is plotted against the observations. The lines present the conditional quantile functions for the .1, .2, \dots , .9 quantiles. Target values equal to 50 correspond to right-censored observations; this was ignored in all analyses as quantile regression forests were not able to deal with right-censoring.






Figure 20: Boston Housing Regression Problem. Average negative log-likelihood for four transformation models computed on 100 validation samples of size N from five data generating processes corresponding to the five panels. The data generating processes were estimated on the original Boston Housing data. For better display, 23 extreme outliers were removed from the plot.

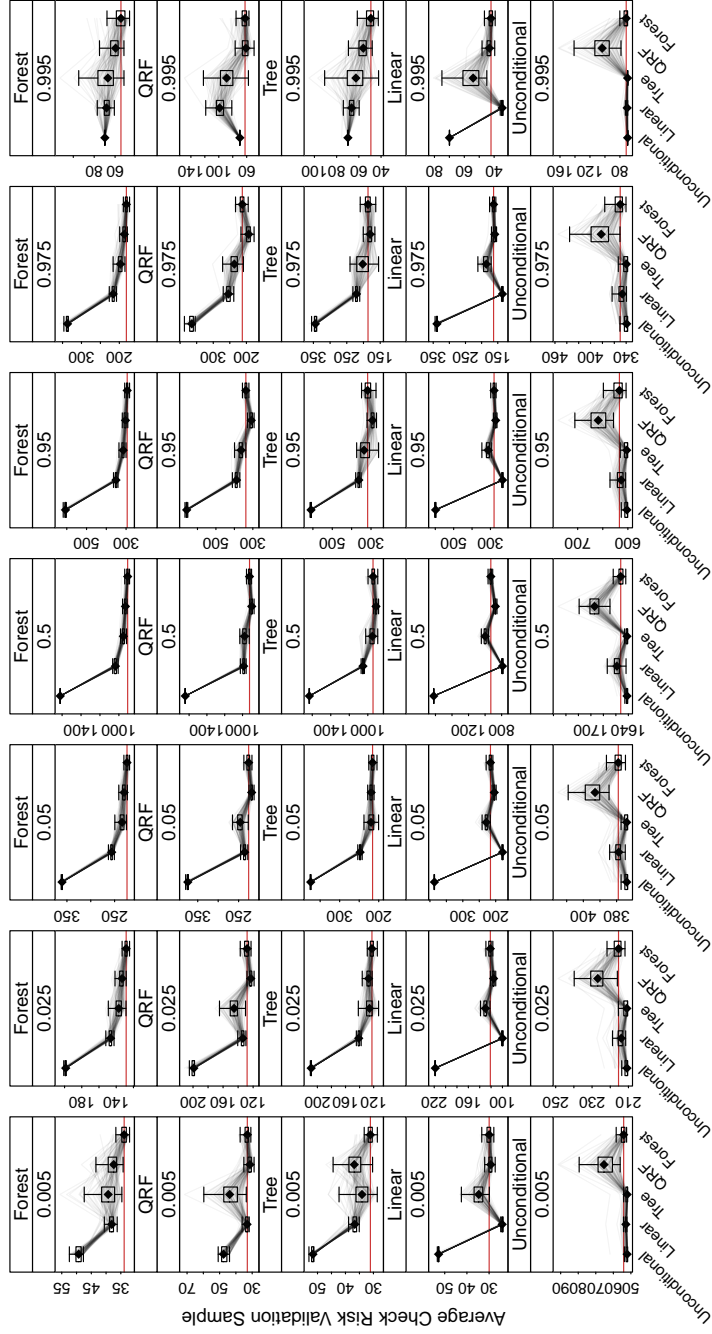


Figure 21: Boston Housing Regression Problem. Average check risk for varying quantiles (in columns) for four transformation models and quantile regression forests (QRF) computed on 100 validation samples of size N from five data generating processes corresponding to rows. The data generating processes were estimated on the original Boston Housing data.